

EXHIBIT G

Design of an Integrated Services *Packet* Network

JONATHAN S. TURNER, MEMBER, IEEE

Abstract—The Integrated Services Digital Network (ISDN) has been proposed as a way of providing integrated voice and data communications services on a universal or near-universal basis. In this paper, I argue that the evolutionary approach inherent in current ISDN proposals is unlikely to provide an effective long-term solution and I advocate a more revolutionary approach, based on the use of advanced packet-switching technology. The bulk of this paper is devoted to a detailed description of an Integrated Services *Packet* Network (ISP), which I offer as an alternative to current ISDN proposals.

An integrated voice and data packet communications system has several advantages over existing methods.

- It uses a common set of switching and transmission facilities for both voice and data communication. This is less costly than current systems that use separate mechanisms.
- It allows voice communication to be done using less than 25 percent of the bandwidth currently needed, without sacrificing quality. This allows major savings in long-distance transmission costs. It also allows customers to carry on two or three simultaneous voice conversations along with a substantial amount of data traffic over a standard copper loop.
- It provides much higher performance data communication and at lower cost than current systems. This is largely due to the integration of data communication with voice, which allows one to take advantage of the economies of scale possible in the large systems needed for a national telephone network.

The system is based on high-performance packet switches which are large and fast enough to effectively support both voice and data communication on a large scale. Each packet switch can have a raw throughput of up to 1.5 Gbits/s, allowing it to support as many as 50 000 simultaneous voice conversations using a 32 kbit/s voice encoding scheme. The one-way cross-network delay in a worst-case connection in a national network in the U.S. can be limited to about 150 ms. The key elements of the design are as follows.

- The use of high-speed digital transmission facilities (1.5 Mbits/s) with excellent error performance.
- Simple link protocols. In particular, there is no flow

Manuscript received April 15, 1986. This work was supported by a Bell Communications Research Grant. This paper was presented at the ACM 9th Data Communications Symposium, Whistler, B.C., Canada, September, 1986.

The author is with the Department of Computer Science, Washington University, St. Louis, MO 63130.
IEEE Log Number 8609694.

control or error correction done at the link level, which eliminates the need for state information in the link level protocol processors.

- A predominantly connection-oriented service. This allows the routing of most packets to be handled by a very simple method and facilitates bandwidth allocation and overload control. A connectionless (datagram) service can also be supported.
- Hardware implementation of basic switching and protocol functions. Switching is done using a large self-routing network containing roughly 1300 custom VLSI chips. All per-packet protocol functions are handled by protocol processors (one for each link) consisting of one custom controller chip plus one large memory chip.
- Implementation of higher level protocol functions (including error correction and flow control) on an end-to-end and application-dependent basis.

The governing philosophy behind the design is that the communications network should provide transport of information at the highest possible level of performance, but nothing else. The network achieves generality, not by providing every service a user might conceivably require, but by providing only those services that every user requires. The system can be implemented with currently available technology at a cost that is comparable to that of conventional telephone switching systems.

THE TROUBLE WITH ISDN

The Integrated Services Digital Network has been heralded as the mechanism that will usher in the Information Age. ISDN, it is said, will facilitate the development of new communications services, including a wide range of data services and maybe even video. It will allow such services to be implemented on a large scale at a cost most customers can afford, and spur the transformation to a "postindustrial society."

While I share the long-range goals of ISDN proponents, I do not believe that current ISDN plans can take us very far. The fundamental problem is that current plans (see [1], for example) implicitly assume a network model based on circuit-switched voice and a combination of circuit and packet-switched data. This assumption is evident in current standardization efforts which focus on a transmission format for digital subscriber lines that provides two 64 kbit/s circuit-switched channels plus a 16 kbit/s packet-switched channel. While this plan can provide a limited data communications capability (it is certainly a vast improvement over the current situation), it is too inflexible to satisfy long-term needs. What happens, for ex-

ample, to an application that requires a 75 kbit/s channel? Do we implement it using the two 64 kbit/s channels or one of them and the packet-switched channel? Neither option is particularly attractive.

The source of the trouble is the reliance on circuit switching, which requires that the available bandwidth be divided up into fixed-size channels. The channel size is a permanent feature which cannot be changed easily, if at all. There is no reason to think that 64 kbit/s is a particularly useful channel size. The only reason for choosing it is that current telephone switching systems are based on that size. In fact, the driving force behind current ISDN plans is to preserve the investment in existing equipment while "evolving" to a more flexible network.

Unfortunately, it will not work because it is inherently a hybrid approach. The only thing integrated about ISDN is its name. Two (or more) switching networks are required to support ISDN as it is currently envisioned. Manufacturers may talk about their integrated architectures for ISDN, but what they mean is that they plan to put a packet switch and a circuit switch in the same box and call it an integrated system. This is not their fault—what else can they do? Packet switching and circuit switching are very different communications methods. They require different switching and transmission facilities and all the proposed schemes for combining them are little more than packaging.

What to do then? Is a hybrid network really necessary or is there an integrated solution that can satisfy the needs of both voice and data and remain flexible enough to meet satisfy new requirements as they arise? I claim that there is such a solution, but it requires abandoning circuit switching and moving to new network designs based on packet switching.

WHY PACKET SWITCHING?

Here are three reasons that make packet switching an attractive method for providing integrated voice and data services.

- *Adaptability to Changing Traffic:* Packet switching naturally provides the user with exactly the bandwidth required. As new services are developed with different bandwidth requirements, packet-switching systems can adapt to the changing conditions easily. Circuit-switching systems cannot.

- *Integrated Internal Architecture:* As outlined above, current ISDN plans require separate switching networks for different types of information. Packet switching can provide both an integrated customer interface and a single network solution for a wide range of communications needs, leading to substantial cost savings in switching systems and system administration.

- *Transmission Efficiency:* Many data services are characterized by bursty communications patterns which make poor use of conventional circuit-switched facilities. For example, interactive data users typically use only a few percent of the bandwidth available to them. Although

it is less widely recognized, voice is also bursty. In the average telephone conversation, less than 40 percent of the available bandwidth is actually used. Packet switching can exploit this burstiness to double the number of conversations that can share a single transmission facility.

Given all these advantages, why has packet switching not been used extensively for voice? To answer this, we must take a closer look at the technical specifications and costs of conventional circuit switches and packet switches and see how they compare.

Circuit switching has been the technology of choice in the telephone network since its origin about 100 years ago. During that period, circuit switches have evolved from manually operated switchboards to small but automatic step-by-step switches to larger panel switches to the computer-controlled electronic switching systems that currently dominate the scene. The current systems are very large—local switches can provide service to over 100 000 customers; large toll switches support as many as 50 000 simultaneous voice calls corresponding to a raw bandwidth of 6 Gbit/s. The network as a whole is also very large. There are approximately 10^8 telephones in the United States at present and over 10 000 local switching offices. The performance of the switching systems is quite impressive—information passing through a modern digital switch is typically delayed no more than a few milliseconds, and new connections can be established in a fraction of a second. Equally impressive is the cost—while there is considerable variation, per-line equipment costs for modern digital telephone systems is typically in the \$100–200 range.

Just as circuit switching has long dominated the telephone network, so has packet switching dominated the data communications scene, principally because of the advantages cited earlier. Packet switching is exemplified by the ARPANET, which was the first major example of a large-scale data communications network. In the ARPANET, the endpoints of the communication are typically large time-shared computers, called hosts. Communication is provided by packet switches, each of which typically connects to a few hosts and several other packet switches. There are currently about 300 hosts in the ARPANET and 100 packet switches. The packet switches are implemented using general-purpose computers (usually minicomputers), although more recent versions have been supplemented with front-end communications processors to reduce the load on the main processor. The transmission facilities used by the ARPANET include low-speed modem connections (1–10 kbit/s) and higher speed digital facilities (56 kbit/s). The throughput of the packet switches is generally under 1 Mbit/s and delays can be substantial (50–100 ms per switch). The cost of packet switching as provided by the ARPANET is quite high since each packet switch supports a small number of hosts. Commercial data networks do better, but there remains a large gap between per-host costs in data networks and per-line costs in the telephone network.

This comparison explains the conventional wisdom that

packet switching is poorly suited to the needs of telephony and the resulting conclusion that an ISDN implementation must include a circuit-switching component to support voice. The fallacy in the conventional wisdom is that the disadvantages attributed to packet switching by this comparison are not due to any inherent properties, but are side effects of the conventional implementations. The requirements and design constraints that shaped the development of the ARPANET and commercial data networks were completely different from those that shaped the telephone network. The scale, the performance requirements, and the cost sensitivities are all very different. It is because the *needs* differ that the resulting systems differ so in their technical specifications. In the remainder of this paper, I will attempt to correct the widely held, but erroneous view that packet switching is poorly suited to the needs of voice by describing a system capable of supporting voice and data communication on a large scale and at a cost that is competitive with conventional telephone systems.

ISPN ARCHITECTURE

Let us begin by deciding what general properties a packet-switching system must have if it is to be a suitable vehicle for providing voice and data communication on a large scale. First and most obviously, it must be big—comparable in raw bandwidth to conventional telephone systems. Second, it must be fast—long end-to-end delays are annoying in voice connections, and hence unacceptable. While opinions differ on the exact amount of delay that can be tolerated, most experts would agree that 100 ms is acceptable, while more than 500 ms is not. Third, it must be inexpensive. Any system that seeks to replace circuit switching must be able to provide voice services at a competitive cost. It is not enough to offer a better product at a higher cost—for a system to succeed on a large scale, it cannot significantly increase costs for basic services.

How can we achieve the requisite scale, performance, and economy in a packet-switching network? Well, the telephone network is one place to look for the answer. One of the first things one notices is that the telephone network makes extensive use of high-bandwidth digital transmission facilities. Modern digital switching systems are designed to interface directly to 1.5 Mbit/s transmission facilities carrying 24 voice channels in a 64 kbit/s format. Interfacing costs are cheap since they are shared by the 24 channels and the bit error rates are excellent. Conventional packet-switching systems, on the other hand, interface to much smaller channels, and must be prepared to cope with the much higher error rates present in modem connections. A second thing one notices about the telephone network is the amount of special-purpose hardware used to provide the switching function. While telephone systems contain general-purpose computers, their function is primarily connection establishment. The

computers do not move the bits. That function is handled by large specialized networks. In contrast, most conventional packet switches include a general-purpose computer that must do some processing on every packet and can quickly become a bottleneck.

These observations suggest that a high-performance packet switch should 1) interface directly to high-speed digital transmission facilities and 2) use special-purpose hardware to perform all per-packet processing. Now, those familiar with link level protocols used in conventional packet networks may recognize this as a challenging task. Typical protocols are quite complex, requiring extensive state information in the protocol processors and complicated error recovery procedures. This complexity almost demands a programmable processor (a hardware implementation would never be completely debugged), but a microprocessor-based implementation is unlikely to be fast enough to keep up with a 1.5 Mbit/s link, and also may be too expensive. Fortunately, there is a way out—simplify the protocol. Current packet-switching protocols were designed for hostile environments—the low speeds and high error rates typical of modem connections. In a network composed entirely of high-speed digital facilities, much of the complexity of typical protocols can be eliminated from the lower protocol layers. In particular, error correction and flow control can be taken out of the link layer and provided on an end-to-end basis rather than a link-by-link basis. This eliminates the need for extensive state information at the ends of each link and the synchronization and recovery procedures required to maintain it, and in turn makes possible the construction of inexpensive, high-performance protocol processors. In addition, if these functions are provided on an end-to-end basis, they can be provided selectively. Hence, delay-sensitive applications like voice can avoid the performance penalties they can cause.

Returning to the question of scale, how large must our packet switches be? In order to compete with conventional telephone switches, they should be capable of supporting at least 50 000 simultaneous voice conversations. How many conversations can a 1.5 Mbit/s link carry? When operated in a circuit-switching mode using the 64 kbit/s digital encoding method currently employed, the answer is 24. When operated in a packet-switching mode, that number jumps to about 50. One can obtain another factor of two improvement by using newer voice encoding methods. 32 kbit/s adaptive differential pulse code modulation (ADPCM) is a good choice since it provides comparable quality to the method currently used, but requires only half the bandwidth. This leads to a figure of 100 voice channels per 1.5 Mbit/s link, implying that our packet switch must terminate 1000 full-duplex links if it is to support 50 000 simultaneous voice conversations. This, in turn, means that our switch must include some mechanism capable of receiving over two million packets per second, and sending each one out on the appropriate outgoing link.

NETWORK OVERVIEW

Based on the discussion in the previous section, we can begin to develop an architecture for the ISPN. The major components are shown in Fig. 1. The packet switches (PS) each terminate up to 1000 high-speed links (HSL). Using 32 kbit/s ADPCM coding for voice, they can support over 50 000 simultaneous voice conversations. Residential customers connect to the network over medium-speed links (MSL), which operate at 100 kbit/s. High-speed access can be provided to business customers. The customer premises interface (CPI) can take a variety of forms, depending on the kind of service required. At the low end of the spectrum would be a simple controller providing service for a single telephone and implemented using a single-chip 8 bit microcomputer. Customers wanting several phones and data communication would require a more complex controller. Businesses would typically have an interface to a private branch exchange (PBX) or local area network (LAN). The network interfaces (NI) provide concentration, accounting data collection, and network protection. A configuration designed for residential customers would support about 500 customers and would be connected to a local packet switch by four HSL's. Thus, one PS could have as many as 125 NI's and support over 60 000 customers.

NI's can also be designed to provide a conventional line interface as supported by current telephone systems. Even higher concentration ratios can be supported in this case—up to 2000 customers could be supported on a single NI connected to its host PS by four HSL's. While this configuration does not exploit the advanced capabilities of the system, it does provide a mechanism for easing the transition from a circuit-oriented network to a packet-oriented one. In a similar fashion, NI's can be designed to interface to the current telephone network. In this case, the NI would interface directly to up to 960 digital trunks.

The network provides two communications services.

- *Point-to-Point Channels*: These are two-way channels joining pairs of customers. The customer establishes a channel by sending a connection request message to the network specifying the destination and the average bandwidth needed. If the network accepts the connection, it in effect guarantees that the customer can expect to have the requested bandwidth available. The network may refuse to accept the connection if there are not adequate resources available. The connections do not provide perfectly reliable information transport. In particular, the network does not provide mechanisms for error correction and flow control. The network can provide connections at any speed up to 1.5 Mbit/s (although larger connections are more likely to be blocked). The bandwidth requirement may be asymmetric—that is, it may depend on the direction of transmission. No distinction is made between voice and data connections. A voice connection is simply one with an average bandwidth of about 12 kbit/s.

- *Datagrams*: These are individually addressed packets, not associated with a preestablished connection. The

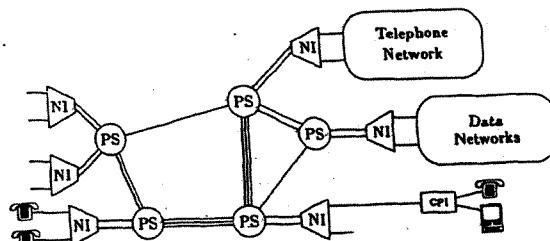


Fig. 1. Network architecture.

network makes an effort to deliver them, but does not guarantee delivery.

The communications protocols are divided into three levels—the *link level*, the *network level*, and the *customer level*. The interrelationships among the different levels are indicated in Fig. 2. There is a single link level protocol, two network level protocols (one for connection-based communication and one for datagrams), and a variety of customer level protocols. The customer level protocols are not discussed in detail here, but the intention is that these protocols would be application-dependent and built on top of one of the two network level protocols. One of these would be a telephony protocol. Another might be an internet protocol such as the DARPA IP to facilitate communication among different data networks. Still another might be a connection-oriented internet protocol.

The link level protocol provides frame delimiting, link transparency, error detection, packet timing, and congestion control, but not error correction. There are four fields used by the link protocol, the frame type field (FTYP), the priority field (PRI), the time stamp field (TS), and the frame check field (FC). The FTYP field identifies the frame as a test frame, a datagram, or a frame belonging to a connection. The priority field (PRI) contains a customer-specified priority. The network preferentially discards low-priority frames to alleviate short-term overload conditions in the network. The network uses the time stamp field (TS) to record the delay encountered by the packet as it crosses the network. More precisely, it records the number of milliseconds the packet spent in each PS and NI it passed through (see [11]). This is important for applications such as voice that are sensitive to delay variations and need a mechanism to remove the timing "jitter" that packet networks can introduce. It can also be used in distributed programming applications for clock synchronization. The frame check field (FC) uses a 16 bit cyclic redundancy code to detect errors in the frame. Frames with errors are simply discarded.

There are two protocols at the network level, one for datagrams and one for connections. Connection packets contain a 4 bit packet type field (PTYP) and a 12 bit logical channel number field (LCN). The packet type field (PTYP) identifies each packet as either a data packet or a control packet and contains a congestion control subfield, used to inform the NI's and customers of internal network congestion. They also contain an information field, which can have any length up to 144 bytes. Control packets are used to establish and control connections and contain a

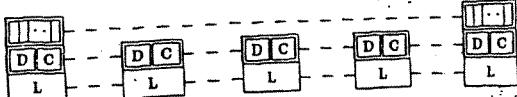


Fig. 2. Protocol structure.

control function field (CF) and a supplementary information field (SI).

PACKET SWITCH DESIGN

The network is built using large high-performance packet-switching systems, each terminating up to 1023 HSL's. The structure of such a packet switch is illustrated in Fig. 3, which shows a small version with 15 HSL's.

The system is controlled by a control processor (CP) which performs all connection control functions, plus administrative and maintenance functions. The CP is a large general-purpose computer. Its role is analogous to that of the control processor in large telephone switching systems such as the No. 4 ESS [2].

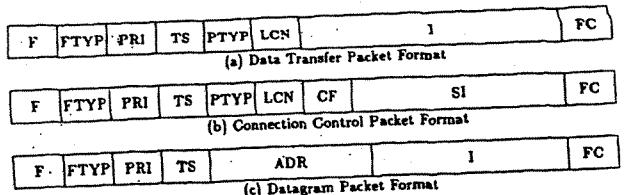
Each HSL is terminated by a packet processor (PP), which performs the link level protocol for all packets and the network level protocol for data transfer packets. It also forwards connection control packets to the CP and datagram packets to the datagram routers (DR).

The heart of the switch is the switch fabric (SF) which consists of a large binary routing network. The important property of such networks is that the path each packet takes through the network is determined by successive bits of its destination address. The figure shows paths from two different SF input ports to output port 1011. Note that at the first stage, the packets are routed out the lower port of the nodes (corresponding to the first "1" bit of the destination address), at the second stage they are routed out the upper port (corresponding to the "0" bit), and in the third and fourth stages they are routed out the lower ports. The self-routing property is shared by a variety of interconnection patterns, including the so-called delta, shuffle-exchange, and banyan networks (see [5]). The PS uses a ten-stage binary routing network with 1024 ports.

The datagram routers (not shown) are special-purpose devices used to route datagrams. The number of DR's can be engineered to suit the traffic. Each occupies a port on the SF, replacing one PP. Since most of the traffic is expected to be connection-oriented, the number of DR's required should be modest.

PACKET PROCESSING

When a packet is received by a PP, it is placed in a buffer with several additional header fields added. The destination field (D) identifies the destination port on the SF. The source field (S) identifies the port where the packet arrived. The length field (LNG) gives the packet length in bytes. The switch packet type field (SPTYP) is used to identify various packet types within the PS. The arrival time field (AT) gives the time at which the packet arrived at the PS (this is used for processing the TS field).



F: Flag (1 byte)
 FTYP: Frame type (4 bits)
 PRI: Priority (4 bits)
 TS: Time stamp (1 byte)
 PTYP: Packet type (4 bits)
 LCN: Logical channel number (12 bits)
 CF: Control function (1 byte)
 SI: Supplementary information
 ADR: Address (8 bytes)
 I: Information field
 FC: Frame check sequence (2 bytes)

Fig. 3. Packet formats.

For data transfer packets, the destination port is determined by the PP using the packet's LCN field and the PP's logical channel translation table (LCXT). Each entry in the LCXT contains an outgoing port number and a new LCN. The outgoing port is placed in the D field of the packet and the new LCN goes in the LCN field. The packet is then sent on to the switch fabric, which uses the D field to route the packet to the proper outgoing port as described earlier. When a packet arrives at the outgoing PP, it is buffered and then transmitted on the HSL with the extra header information stripped off.

The contents of the LCXT's is controlled by the CP, which can read and write the LCXT using special control packets sent to the PP's through the SF. Thus, the connection establishment process includes the sending of messages from the CP to the two PP's selected for the connection, updating their LCXT's appropriately.

SWITCH FABRIC

The basic operation of the switch fabric has already been described. The SF's highly regular and parallel structure allows the construction of very large switches, without the bottlenecks that can arise in bus or ring-based interconnection networks.

The nodes of the SF operate as miniature packet switches. Each node has a buffer at each input port capable of holding one maximum length packet. The data paths joining the nodes are bit serial and operate at 12 Mbits/s. This gives the SF an 8:1 speed advantage over the external HSL's. Thus, if all the external links are operated at an occupancy of 85 percent, the internal links will have an average occupancy of less than 11 percent. (This is assuming just one switch plane active. When both are active, the average occupancy is less than 6 percent.) There is also an upstream control lead joining each pair of adjacent nodes. This is used to implement a simple hardware flow control mechanism, which prevents buffer overflows within the SF.

A more detailed look at the switch node appears in Fig. 4. It consists of two input controllers (IC) and two output controllers (OC). The IC's contain a buffer large enough to hold one packet and a controller implemented as a state machine. The OC's are simple state machines that arbitrate requests for their ports.

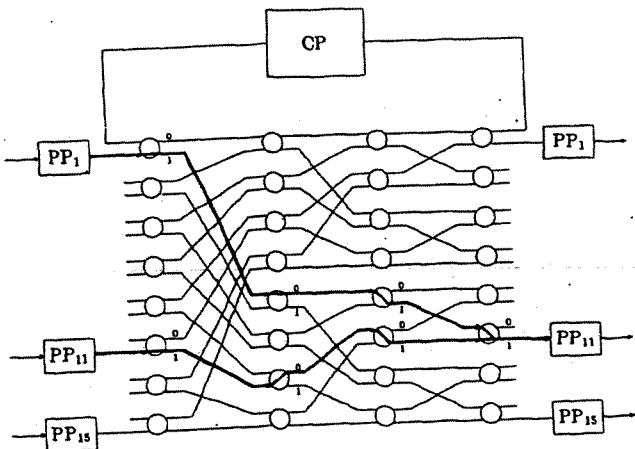


Fig. 4. Packet switch structure.

When a packet is received, the IC determines the proper outgoing port by examining the appropriate bit of the packet's destination field, then requests permission to use that port. If the desired port is immediately available, the packet is sent to it directly, bypassing the buffer. Hence, a packet can pass through a switch node after experiencing a delay of just a few bit times. In fact, this is the normal case due to the relatively low occupancy of the internal data paths.

If the desired port is not available, the packet is shifted into the buffer. As soon as the desired port becomes available, the packet is sent out—even if not all of the packet has been received. If the port is still not available when the end of the packet is received, the IC holds its *grant* lead low to prevent the arrival of new packets. The *grant* lead is reasserted as soon as the desired link becomes available, allowing a new packet to enter the buffer, while another is leaving.

PACKET PROCESSORS

The structure of the packet processors (PP) is shown in Fig. 5. It is organized around a 16 kbyte RAM, which contains four packet buffers and the logical channel translation table (LCXT). The principal buffers are the receive buffer (RCB) used for packets received from the HSL on their way to the switch, and the transmit buffer (XMB) for packets going from the switch to the HSL. The link test buffer (LTB) and the switch test buffer (STB) are small buffers that provide loop-back paths for testing the HSL and switch, respectively.

Access to the memory is provided through the address controller (ADC), which contains read and write pointers for the buffers and arbitrates memory access.

The receive circuit (RCV) receives the incoming packets from the HSL, removes the flag field, discards packets with errors, adds the extra header fields, initializes the length (LNG) and arrival time (AT) fields, converts from bit serial to 8 bit parallel format, and writes the packet to the RCB through the ADC.

The output circuit (OUT) takes packets from the RCB,

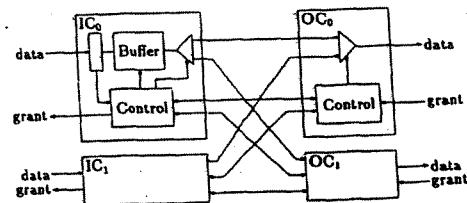


Fig. 5. Switch node.

performs the logical channel translation described above, and sends the packets onto the switch in bit serial format. The input circuit (IN) takes packets from the switch and writes them to the XMB.

The transmit circuit (XMIT) takes packets from the XMB, performs the time stamp calculation, strips the extra header information, adds the flag field, and transmits the packets on the HSL.

The switch interface (SI) connects to the duplicated switch planes, normally routing packets to and receiving packets from the active plane. The SI can also send and receive packets from the standby plane—this is used for testing.

The PP can be implemented using two chips—one custom controller chip and one memory chip. The simplicity of the protocols makes this possible. There is no need to buffer unacknowledged frames that may need to be retransmitted as in conventional link level protocols that perform error correction and flow control. Similarly, there is no need for the recovery and synchronization procedures that are required by such protocols.

One problem with binary routing networks is that they can become congested in the presence of certain traffic patterns. This is illustrated in Fig. 6, which shows a traffic pattern corresponding to several *communities of interest*. In this pattern, all traffic entering the first four inputs is destined for the first four outputs, all traffic entering the second group of four inputs is destined for the second group of four outputs, and so forth. Note that with this pattern, only one fourth of the links joining the second and third stages are carrying traffic. Thus, if the inputs are heavily loaded, the internal links will be hopelessly overloaded and traffic will back up. In a 1024×1024 network, there are ten stages and the links between the fifth and sixth stages can, in the worst case, be carrying all the traffic on just 32 of the 1024 links.

This problem can be solved by using two networks instead of one. One of these is called the routing network (RN) and is a standard binary routing network. The other network sits in front of the RN and is called the distribution network (DN). It has the same structure as the RN, but instead of routing packets based on their destination address, it attempts to distribute packets evenly across all its output ports. This is done by having each switch node route packets alternately out its two ports. The alternate-port strategy is modified if one or both ports is unavailable. In this case, the first port to become available is used. This approach breaks up any communities of inter-

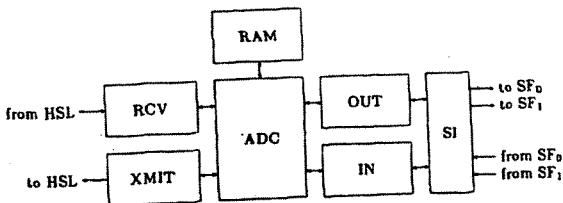


Fig. 6. Packet processor.

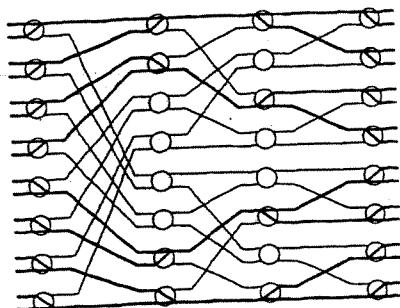


Fig. 7. Congestion in binary routing networks.

est and makes the combination of the DN and RN robust in the face of pathological traffic patterns.

One drawback of the DN is that it doubles the number of stages in the SF, thus roughly doubling the packet delay and the circuit complexity. This loss can be recovered by using larger switch nodes for the RN and DN. Instead of 2×2 , we can use 4×4 nodes. With the larger nodes, we can construct a network with half the number of stages required with 2×2 nodes. This also halves the delay and the circuit complexity.

SUMMARY

This paper has described the design of a high-performance packet-switching network capable of supporting both voice and data communication on a large scale and at a cost to the user comparable to current telephone service. The key features of the design are the use of high-speed digital transmission facilities and simple link level protocols, a predominantly connection-oriented service, hardware implementation of all per-packet functions, and implementation of higher level protocol functions on an end-to-end and application-dependent basis. I claim that the advanced packet-switching technology described here is inherently better suited for the provision of advanced communications services than the circuit-switching technology which underlies current ISDN proposals.

Many aspects of the packet-switching system described here have been patented by AT&T Bell Laboratories. See [14]–[20] for further details.

ACKNOWLEDGMENT

The work described here was carried out between 1981 and 1983 when I was a member of the Technical Staff at Bell Laboratories, Naperville, IL. Many individuals contributed ideas. I particularly want to thank H. Andrews, D. Creed, B. Hoberecht, W. Montgomery, and L. Wyatt

who helped turn the original idea into a workable design. I also want to thank B. Cardwell and E. Nussbaum for recognizing the potential and providing the backing to develop it. Special thanks go to M. Décina, whose refusal to accept half-baked ideas forced me to make sure that mine were always well-cooked.

REFERENCES

- [1] D. J. Aoki, D. H. Florin, S. D. McKenna, G. R. Welsh, and P. E. White, "Digital feature evolution in the SESS™ digital switch," in *Proc. GLOBECOM 83*, 1983, pp. 601–605.
- [2] *Bell Syst. Tech. J.*, vol. 56, Sept. 1977 (Special Issue devoted to the No. 4 ESS).
- [3] M. Décina, "Performance requirements for integrated voice/data networks," *IEEE Trans. Commun.*, vol. COM-30, pp. 2117–2130, Sept. 1982.
- [4] D. M. Dias and J. R. Jump, "Packet switching interconnection networks for modular systems," *Comput.*, vol. 14, pp. 43–54, Dec. 1983.
- [5] T.-Y. Feng, "A survey of interconnection networks," *Comput.*, vol. 14, pp. 12–30, Dec. 1983.
- [6] J. G. Gruber, "Performance requirements for integrated voice/data networks," *IEEE J. Select. Areas Commun.*, vol. SAC-1, pp. 981–1005, Dec. 1983.
- [7] W. L. Hoberecht, "A layered network protocol for packet voice and data integration," *IEEE J. Select. Areas Commun.*, vol. SAC-1, pp. 1006–1013, Dec. 1983.
- [8] Y.-C. Jeng, "Performance analysis of a packet switch based on a single-buffered banyan network," *IEEE J. Select. Areas Commun.*, vol. SAC-1, pp. 1014–1021, Dec. 1983.
- [9] M. Kasahara, H. Shimizu, and M. Imaizumi, "Basic concept of the digital network for the information network system," in *Proc. GLOBECOM 83*, 1983, pp. 969–973.
- [10] P. Kermani and L. Kleinrock, "Virtual cut-through: A new computer communication switching technique," *Comput. Networks*, vol. 3, pp. 267–286, 1979.
- [11] W. A. Montgomery, "Techniques for packet voice synchronization," *IEEE J. Select. Areas Commun.*, vol. SAC-1, pp. 1022–1028, Dec. 1983.
- [12] R. Reitberg, C. Wyman, D. Hunt, M. Hoffman, P. Carvey, B. Hyde, W. Clark, and M. Kraley, "Development of a voice funnel system: Design report," Bolt Beranek and Newman, Rep. 4098, Aug. 1979.
- [13] J. S. Turner and L. F. Wyatt, "A packet network architecture for integrated services," in *Proc. GLOBECOM 83*, 1983, pp. 45–50.
- [14] J. S. Turner, "Packet load monitoring by trunk controllers," U.S. Patent 4 484 326, Nov. 20, 1984.
- [15] —, "Packet switching loop-around network and facilities testing," U.S. Patent 4 486 877, Dec. 4, 1984.
- [16] —, "End-to-end information memory arrangement in a line controller," U.S. Patent 4 488 288, Dec. 11, 1984.
- [17] —, "Interface facility for a packet switching system," U.S. Patent 4 488 289, Dec. 11, 1984.
- [18] —, "Packet error rate measurements by distributed controllers," U.S. Patent 4 490 817, Dec. 25, 1984.
- [19] —, "Fast packet switch," U.S. Patent 4 491 945, Jan. 1, 1985.
- [20] —, "Fast packet switching system," U.S. Patent 4 494 230, Jan. 15, 1985.
- [21] C. J. Weinstein and J. W. Forgie, "Experience with speech communication in packet networks," *IEEE J. Select. Areas Commun.*, vol. SAC-1, pp. 1022–1028, Dec. 1983.



Jonathan S. Turner (M'77) received the B.S.C.S. and B.S.E.E. degrees from Washington University, St. Louis, MO, in 1977, and the M.S. and Ph.D. degrees in computer science from Northwestern University, Evanston, IL, in 1979 and 1981.

From 1977 to 1983 he worked for Bell Laboratories, in Naperville, IL, first as a member of the Technical Staff and later as a Consultant. His work there included the development of maintenance software and design of system architectures

for telephone switching systems. From 1981 to 1983 he was the principal system architect for the fast packet switching project, an applied research project which established the feasibility of integrated voice and data communication using packet-switching technology. He has been awarded seven patents for this work and has several others pending. He is now an Associate Professor of Computer Science at Washington University, where he con-

tinues his research on high-performance communications systems. His research interests also include the study of algorithms and computational complexity, with particular interest in the probable performance of heuristic algorithms for NP-complete problems.

Dr. Turner is a member of the Association for Computing Machinery and the Society for Industrial and Applied Mathematics.